

AN ALGORITHM FOR FINDING POSITIVE SOLUTIONS TO POLYNOMIAL EQUATIONS

DUSTIN CARTWRIGHT

ABSTRACT. We present a numerical algorithm for finding real non-negative solutions to polynomial equations. Our methods are based on the expectation maximization and iterative proportional fitting algorithms, which are used in statistics to find maximum likelihood parameters for certain classes of statistical models. Since our algorithm works by iteratively improving an approximate solution, we find approximate solutions in the cases when there are no exact solutions, such as overconstrained systems.

We present an iterative numerical method for finding non-negative solutions and approximate solutions to systems of polynomial equations. We require two assumptions about our system of equations. First, for each equation, all the coefficients other than the constant term must be non-negative. Second, there is a technical assumption on the exponents, described at the beginning of Section 1, which, for example, is satisfied if all non-constant terms have the same total degree. In Section 3, there is a discussion of the range of possible systems which can arise under these hypotheses.

Because of the assumption on signs, we can write our system of equations as,

$$(1) \quad \sum_{\alpha \in S} a_{i\alpha} x^\alpha = b_i \quad \text{for } i = 1, \dots, \ell,$$

where the coefficients $a_{i\alpha}$ are non-negative and the b_i are positive, and $S \subset \mathbb{R}_{\geq 0}^n$ is a finite set of possibly non-integer multi-indices. Our algorithm works by iteratively decreasing the generalized Kullback-Leibler divergence of the left-hand side and right-hand side of (1). The generalized Kullback-Leibler divergence of two positive vectors a and b is defined to be

$$(2) \quad D(a||b) := \sum_i \left(b_i \log \left(\frac{b_i}{a_i} \right) - b_i + a_i \right).$$

The standard Kullback-Leibler consists only of the first term and is defined only for probability distributions, i.e. the sum of each vector is 1. The last two terms are necessary so that the generalized divergence has, for arbitrary positive vectors a and b , the property of being non-negative and zero exactly when a and b are equal (Proposition 3).

Our algorithm converges to local minima of the Kullback-Leibler divergence, including exact solutions to the system (1). In order to find multiple local minima, we can repeat the algorithm for randomly chosen starting points. For finding approximate solutions, this may be sufficient. However, there are no guarantees of completeness for the exact solutions obtained in this way. Nonetheless, we hope that in certain situations, our algorithm will find applications both for finding exact and approximate solutions.

Lee and Seung applied the EM algorithm to the problem of non-negative matrix factorization [6]. They introduced the generalized Kullback-Leibler divergence in (2) and used it to find approximate non-negative matrix factorizations. Since the product of two matrices can be expressed by polynomials in the entries of the matrices, matrix-factorization is a special case of the equations in (1).

For finding exact solutions to arbitrary systems of polynomials, there are a variety of approaches which find all complex or all real solutions. Homotopy continuation methods find all complex roots of a system of equations [8]. Even to find only positive roots, these two methods find all complex or all real solutions, respectively. Lasserre, Laurent, and Rostalski have applied semidefinite programming to find all real solutions to a system of equations and a slight modification of their algorithm will find all positive real solutions [4, 5]. Nonetheless, neither of these methods has any notion of approximate solutions.

For directly finding only positive real solutions, Bates and Sottile have proposed an algorithm based on fewnomials bounds on the number of real solutions [1]. However, their method is only effective when the number of monomials (the set S in our notation) is only slightly more than the number of variables. Our method only makes weak assumptions on the set of monomials, but stronger assumptions on the coefficients.

Our inspiration comes from tools for maximum likelihood estimation in statistics. Parameters which maximize the likelihood are exactly the parameters such that the model probabilities are closest to the empirical probabilities, in the sense of minimizing Kullback-Leibler divergence. Expectation-Maximization [7, Sec. 1.3] and Iterative Proportional Fitting [3] are well-known iterative methods for maximum likelihood estimation. We re-interpret these algorithms as methods for approximating solutions to polynomial equations, in which case their applicability can be somewhat generalized.

The impetus behind the work in this paper was the need to find approximate positive solutions to systems of bilinear equations in [2]. In this case the variables represented expression levels, which only made sense as positive parameters. Moreover, in order to accomodate noise in the data, there were more equations than variables, so it was necessary to find approximate solutions. Thus, the algorithm described in this paper was the most appropriate tool. Here, we generalize beyond bilinear equations and present proofs.

An implementation of our algorithm in the C programming language is freely available at <http://math.berkeley.edu/~dustin/pos/>.

In Section 1, we describe the algorithm and the connection to maximum likelihood estimation. In Section 2, we prove that the necessary convergence for our algorithm. Finally, in Section 3, we show that even with our restrictions on the form of the equations, there can be exponentially positive real solutions.

1. ALGORITHM

We make the assumption that we have an $s \times n$ non-negative matrix g , with no column identically zero, and positive real numbers d_j for $1 \leq j \leq s$ such that for each $\alpha \in S$ and each $j \leq s$, $\sum_{i=1}^n g_{ji}\alpha_i$ is either 0 or d_j . For example, if all the monomials x^α have the same total degree d_1 , we can take $s = 1$ and $g_{1i} = 1$ for all i . The other case of particular interest is multilinear systems of equations, in which each α_i is at most one. In this case the variables can be partitioned into sets

such that the equations are linear in each set of variables, so we can take $d_j = 1$ for all j . Note that because d_j is in the denominator in (4), convergence is fastest when the d_j are small, such as in the multilinear case. We also note that, for an arbitrary set of exponents S , there may not exist such a g .

The algorithm begins with a randomly chosen starting vector and iteratively improves it through two nested iterations:

- Initialize x with n randomly chosen positive real numbers.
- Loop until the vector x stabilizes:
 - For all $\alpha \in S$, compute

$$(3) \quad w_\alpha := \sum_i b_i \frac{a_{i\alpha} x^\alpha}{\sum_\beta a_{i\beta} x^\beta}.$$

- Loop until the vector x stabilizes:
 - Loop for j from 1 to s :
 - Simultaneously update all entries of x :

$$(4) \quad x_i \leftarrow x_i \left(\frac{\sum_\alpha \alpha_i g_{ji} w_\alpha}{\sum_\alpha \alpha_i g_{ji} a_\alpha x^\alpha} \right)^{g_{ji}/d_j} \quad \text{where } a_\alpha = \sum_i a_{i\alpha}.$$

Because there is no subtraction, it is clear that the entries of x remain positive throughout this algorithm.

Our method is inspired by interpreting the equations in (1) as a maximum likelihood problem for a statistical model and applying the well-known methods of Expectation-Maximization (EM) and Iterative Proportional Fitting (IPF). Here, we assume that all the monomials x^α have the same total degree. Our statistical model is that a hidden process generates an integer $i \leq \ell$ and an exponent vector α with joint probability $a_{i\alpha} x^\alpha$. The vector x contains n positive parameters for the model, restricted such that the total probability $\sum_{i,\alpha} a_{i\alpha} x^\alpha$ is 1. The empirical data consists of repeated observations of the integer i , but not the exponent α , and b_i is the proportion of observations of i . In this situation, the vector x which minimizes the divergence of (1) is the maximum likelihood parameters for the empirical distribution b_i . The inner loop of the algorithm consists of using IPF to solve the log-linear hidden model and the outer loop consists of using EM to estimate the distribution on the hidden states.

2. PROOF OF CONVERGENCE

In this section we prove our main theorem:

Theorem 1. *The Kullback-Leibler divergence*

$$(5) \quad \sum_{i=1}^{\ell} D \left(\sum_{\alpha \in S} a_{i\alpha} x^\alpha \parallel b_i \right).$$

is weakly decreasing during the algorithm in Section 1. Moreover, assuming that the set S contains a multiple of each unit vector e_i , i.e. some power of each x_i appears in the system of equations, then the vector x converges to a critical point of the function (5) or the boundary of the positive orthant.

Remark 2. *The condition on S is necessary to ensure that the vector x remains bounded during the algorithm.*

We begin by establishing several basic properties of the generalized Kullback-Leibler divergence in Proposition 3 and Lemmas 4 and 5. Note that $D(a\|b) = \sum_{i=1}^m D(a_i\|b_i)$ and thus, we will check these basic properties by reducing to the case where a and b are scalars.

The proof of Theorem 1 itself is divided into two parts, corresponding to the two nested iterative loops. The first step is to prove that the updates (4) in the inner loop converge a local minimum of the divergence $D(a_\alpha x^\alpha\|w_\alpha)$. The second step is to show that this implies that the outer loop strictly decreases the divergence function (5) exact at a critical point.

Proposition 3. *For a and b vectors of positive real numbers, the divergence $D(a\|b)$ is non-negative with $D(a\|b) = 0$ if and only if $a = b$.*

Proof. It suffices to prove this when a and b are scalars. We let $t = b/a$, and,

$$D(a\|b) = b \log\left(\frac{b}{a}\right) - b + a = a(t \log t - t + 1) = a \int_1^t \log s \, ds.$$

It is clear that $\int_1^t \log s \, ds$ is non-negative and equal to zero if and only if $t = 1$. \square

Lemma 4. *Suppose that a and b are vectors of m positive real numbers. Let t be any positive real number, and then*

$$D(ta\|b) = D(a\|b) + (t-1) \sum_{i=1}^m a_i - \sum_{i=1}^m b_i \log t$$

Proof. Again, we assume that a and b are scalars and then it becomes a straightforward computation. \square

Lemma 5. *If a and b are vectors of m positive real numbers, then we can relate their divergence to the divergence of their sums by*

$$D(a\|b) = D\left(\sum_{i=1}^m a_i \parallel \sum_{i=1}^m b_i\right) + D\left(\frac{\sum_{i=1}^m b_i}{\sum_{i=1}^m a_i} a\|b\right).$$

Proof. We let $A = \sum_{i=1}^m a_i$ and $B = \sum_{i=1}^m b_i$, and apply Lemma 4 to the last term:

$$\begin{aligned} D\left(\frac{B}{A}a\|b\right) &= D(a\|b) + \left(\frac{B}{A} - 1\right)A - B \log \frac{B}{A} \\ &= D(a\|b) - B \log \frac{B}{A} + B - A = D(a\|b) - D(A\|B). \end{aligned}$$

After rearranging, we get the desired expression. \square

Lemma 6. *The update rule (4) weakly decreases the divergence. If*

$$x'_i = x_i \left(\frac{\sum_{\alpha} \alpha_i g_{ji} w_{\alpha}}{\sum_{\alpha} \alpha_i g_{ji} a_{\alpha} x^{\alpha}} \right)^{g_{ji}/d_j},$$

then

(6)

$$\sum_{\alpha} D(a_{\alpha}(x')^{\alpha}\|w_{\alpha}) \leq \sum_{\alpha} D(a_{\alpha}x^{\alpha}\|w_{\alpha}) - \frac{1}{d} \sum_{i=1}^n D\left(\sum_{\alpha} \alpha_i g_{ji} a_{\alpha} x^{\alpha} \parallel \sum_{\alpha} \alpha_i g_{ji} w_{\alpha}\right).$$

Proof. First, since the statement only depends on the j th row of the matrix g , we can assume that g is a row vector and we drop j from future subscripts. Second, we can assume that $d = 1$ by replacing g_i with g_i/d .

Third, we reduce to the case when $g_i = 1$ for all i . We define a new set of exponents $\tilde{\alpha}$ and coefficients $\tilde{a}_{\tilde{\alpha}}$ by $\tilde{\alpha}_i = g_i \alpha_i$ and $\tilde{a}_{\tilde{\alpha}} = a_{\alpha} \prod x_i$, where the product is taken over all indices i such that $g_i = 0$. We take \tilde{x} to be a vector indexed by those i such that $g_i \neq 0$. Then, under the change of coordinates $\tilde{x}_i = x_i^{1/g_i}$, we have $a_{\alpha} x^{\alpha} = \tilde{a}_{\tilde{\alpha}} \tilde{x}^{\tilde{\alpha}}$ and the update rule in (4) is the same for the new system with coefficients $\tilde{a}_{\tilde{\alpha}}$ and exponents $\tilde{\alpha}$. Furthermore, if all entries of $\tilde{\alpha}$ are zero, then $\tilde{x}^{\tilde{\alpha}} = 1$ for all vectors x and so we can drop $\tilde{\alpha}$ from our exponent set. Therefore, for the rest of the proof, we drop the tildes, and assume that $\sum_i \alpha_i = 1$ for all $\alpha \in S$ and $g_i = 1$ for all i , in which case g drops out of the equations.

To prove the desired inequality, we substitute the updated assignment x' into the definition of Kullback-Leibler divergence:

$$\begin{aligned}
 D(a_{\alpha}(x')^{\alpha} \| w_{\alpha}) &= w_{\alpha} \log \left(\frac{w_{\alpha}}{a_{\alpha} x^{\alpha} \prod_{i=1}^n \left(\frac{\sum_{\beta} w_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} \right)^{\alpha_i}} \right) - w_{\alpha} + a_{\alpha}(x')^{\alpha} \\
 &= w_{\alpha} \log \frac{w_{\alpha}}{a_{\alpha} x^{\alpha}} - \sum_{i=1}^n \alpha_i w_{\alpha} \log \frac{\sum_{\beta} \beta_i w_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} - w_{\alpha} + a_{\alpha}(x')^{\alpha} \\
 (7) \quad &= D(a_{\alpha} x^{\alpha} \| w_{\alpha}) - \sum_{i=1}^n \alpha_i w_{\alpha} \log \frac{\sum_{\beta} \beta_i w_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} - a_{\alpha} x^{\alpha} + a_{\alpha}(x')^{\alpha}.
 \end{aligned}$$

On the other hand, let C denote the last term of (6), which we can expand as,

$$\begin{aligned}
 C &= \sum_{i=1}^n D \left(\sum_{\alpha} \alpha_i a_{\alpha} x^{\alpha} \| \sum_{\alpha} \alpha_i w_{\alpha} \right) \\
 &= \sum_{i=1}^n \left(\left(\sum_{\alpha} \alpha_i w_{\alpha} \right) \log \frac{\sum_{\alpha} \alpha_i w_{\alpha}}{\sum_{\alpha} \alpha_i a_{\alpha} x^{\alpha}} - \sum_{\alpha} \alpha_i w_{\alpha} + \sum_{\alpha} \alpha_i a_{\alpha} x^{\alpha} \right) \\
 &= \sum_{i=1}^n \sum_{\alpha} \left(\alpha_i w_{\alpha} \log \frac{\sum_{\beta} \beta_i w_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} - \alpha_i w_{\alpha} + \alpha_i a_{\alpha} x^{\alpha} \right) \\
 (8) \quad &= \sum_{\alpha} \left(\sum_{i=1}^n \alpha_i w_{\alpha} \log \frac{\sum_{\beta} \beta_i w_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} \right) - w_{\alpha} + a_{\alpha} x^{\alpha},
 \end{aligned}$$

where the last step follows from the assumption that $\sum_i \alpha_i = 1$ for all $\alpha \in S$. We take the sum of (7) over all $\alpha \in S$ and add it to (8) to get,

$$(9) \quad \sum_{\alpha} D(a_{\alpha}(x')^{\alpha} \| w_{\alpha}) + C = \sum_{\alpha} D(a_{\alpha} x^{\alpha} \| w_{\alpha}) - \sum_{\alpha} b_{\alpha} + \sum_{\alpha} a_{\alpha}(x')^{\alpha}.$$

Finally, we expand the last term of (8) using the definition of x' and apply the arithmetic-geometric mean inequality,

$$\begin{aligned}
\sum_{\alpha} a_{\alpha} (x')^{\alpha} &= \sum_{\alpha} a_{\alpha} x^{\alpha} \prod_{i=1}^n \left(\frac{\sum_{\beta} \beta_i b_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} \right)^{\alpha_i} \\
&\leq \sum_{\alpha} a_{\alpha} x^{\alpha} \sum_{i=1}^n \alpha_i \frac{\sum_{\beta} \beta_i b_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} \\
&= \sum_{i=1}^n \left(\sum_{\alpha} \alpha_i a_{\alpha} x^{\alpha} \right) \frac{\sum_{\beta} \beta_i b_{\beta}}{\sum_{\beta} \beta_i a_{\beta} x^{\beta}} \\
&= \sum_{i=1}^n \sum_{\beta} \beta_i b_{\beta} = \sum_{\beta} b_{\beta}.
\end{aligned}$$

Together with (9), this gives the desired inequality. \square

Proposition 7. *A positive vector x is a fixed point of the update rule (4) for all $1 \leq j \leq s$ if and only if x is a critical point of the divergence function $\sum_{\alpha} D(a_{\alpha} x^{\alpha} \| w_{\alpha})$.*

Proof. For the update rule to be constant means that the numerator and denominator in (4) are equal, i.e.

$$(10) \quad \sum_{\alpha} \alpha_i g_{ji} a_{\alpha} x^{\alpha} = \sum_{\alpha} \alpha_i g_{ji} w_{\alpha} \quad \text{for all } i \text{ and } j.$$

By our assumption on g , for each i , some g_{ji} is non-zero, so (10) is equivalent to

$$(11) \quad \sum_{\alpha} \alpha_i a_{\alpha} x^{\alpha} = \sum_{\alpha} \alpha_i w_{\alpha} \quad \text{for all } i.$$

On the other hand, we compute the partial derivative

$$\frac{\partial}{\partial x_i} \sum_{\alpha} D(a_{\alpha} x^{\alpha} \| w_{\alpha}) = \sum_{\alpha} -w_{\alpha} \frac{\alpha_i}{x_i} + \alpha_i a_{\alpha} \frac{x^{\alpha}}{x_i}.$$

Since each x_i is assumed to be non-zero, it is clear that all partial derivatives being zero is equivalent to (11). \square

Lemma 8. *If we define w_{α} as in (3), then*

$$\begin{aligned}
\sum_{i=1}^n D \left(\sum_{\alpha} a_{i\alpha} (x')^{\alpha} \| b_i \right) &- \sum_{i=1}^n D \left(\sum_{\alpha} a_{i\alpha} x^{\alpha} \| b_i \right) \\
&\leq \sum_{\alpha} D(a_{\alpha} (x')^{\alpha} \| w_{\alpha}) - \sum_{\alpha} D(a_{\alpha} x^{\alpha} \| w_{\alpha}).
\end{aligned}$$

Moreover, a positive vector x is a fixed point if and only if x is a critical point for the divergence function.

Proof. We consider

$$(12) \quad \sum_{i,\alpha} D(a_{i\alpha} (x')^{\alpha} \| w_{i\alpha}) \quad \text{where } w_{i\alpha} = \frac{b_i a_{i\alpha} x^{\alpha}}{\sum_{\beta} a_{i\beta} x^{\beta}},$$

and apply Lemma 5 in two different ways. First, by applying Lemma 5 to each group of (12) with fixed α , we get

$$\sum_{i,\alpha} D(a_{i\alpha}(x')^\alpha \| w_{i\alpha}) = \sum_{\alpha} D(a_{\alpha}(x')^\alpha \| w_{\alpha}) + \sum_{i,\alpha} D\left(\frac{\sum_j w_{j\alpha}}{\sum_j a_{\alpha}(x')^\alpha} a_{i\alpha}(x')^\alpha \| w_{i\alpha}\right).$$

In the last term, the monomials $(x')^\alpha$ cancel and so it is a constant independent of x' which we denote E . On the other hand, we can apply Lemma 5 to each group in (12) with fixed i ,

$$\sum_{i,\alpha} D(a_{i\alpha}(x')^\alpha \| w_{i\alpha}) = \sum_i D\left(\sum_{\alpha} a_{i\alpha}(x')^\alpha \| b_i\right) + \sum_{i,\alpha} D\left(\frac{b_i a_{i\alpha}(x')^\alpha}{\sum_{\beta} a_{i\beta}(x')^\beta} \| w_{i\alpha}\right).$$

We can combine these equations to get

$$(13) \quad \sum_i D\left(\sum_{\alpha} a_{i\alpha}(x')^\alpha \| b_i\right) = \sum_{\alpha} D(a_{\alpha}(x')^\alpha \| w_{\alpha}) + E - \sum_{i,\alpha} D\left(\frac{b_i a_{i\alpha}(x')^\alpha}{\sum_{\beta} a_{i\beta}(x')^\beta} \| w_{i\alpha}\right).$$

By Proposition 3, the last term of 13 is non-negative, and by the definition of $w_{i\alpha}$, it is zero for $x' = x$. Therefore, any value of x' which decreases the first term compared to x will also decrease the left hand side by at least as much, which is the desired inequality.

In order to prove the statement about the derivative, we consider the derivative of (13) at $x' = x$. Because the last term is minimized at $x' = x$, its derivative is zero, so

$$\left. \frac{\partial}{\partial x'_j} \right|_{x'=x} \sum_i D\left(\sum_{\alpha} a_{i\alpha}(x')^\alpha \| b_i\right) = \left. \frac{\partial}{\partial x'_j} \right|_{x'=x} \sum_{i,\alpha} D(a_{i\alpha}(x')^\alpha \| w_{i\alpha}).$$

By Proposition 7, a positive vector x is a fixed point of the inner loop if and only if these partial derivatives on the right are zero for all indices j , which is the definition of a critical point. \square

Proof of Theorem 1. The Kullback-Leibler divergence $\sum_{\alpha} D(a_{\alpha} x^{\alpha} \| w_{\alpha})$ decreases at each step of the inner loop by Lemma 6. Thus, by Lemma 8, the divergence

$$(14) \quad \sum_{i=1}^n D\left(\sum_{\alpha} a_{i\alpha} x^{\alpha} \| b_i\right)$$

decreases at least as much. However, the divergence (14) is non-negative according to Proposition 3. Therefore, the magnitude of the decreases in divergence must approach zero over the course of the algorithm. By Lemma 6, this means that the quantity C in that theorem must approach zero. By Proposition 3, this means that the quantities in that divergence approach each other. However, up to a power, these are the numerator and denominator of the factor in the update rule (4), so the difference between consecutive vectors x approaches zero.

Thus, we just need to show that x remains bounded. However, since some power of each variable x_i occurs in some equation, as x_i gets large, the divergence for that equation also gets arbitrarily large. Therefore, each x_i must remain bounded, so the vector x must have a limit as the algorithm is iterated. If this limit is in the interior of the positive orthant, then it must be a fixed point. By Lemma 8 and Proposition 7, this fixed point must be a critical point of the divergence (5). \square

3. UNIVERSALITY

Although the restriction on the exponents and especially the positivity of the coefficients seem like strong conditions, such systems can nonetheless be quite complex. In this section, we investigate the breadth of such equations.

Proposition 9. *For any system of ℓ real polynomials in n variables, there exists a system of $\ell + 1$ equations in $n + 1$ variables, in the form (1), such that the positive solutions (x_1, \dots, x_n) to the former system are in bijection with the positive solutions (x'_1, \dots, x'_{n+1}) of the latter, with $x'_i = x_i/x_{n+1}$.*

Proof. We write our system of equations as $\sum_{\alpha \in S} a_{i\alpha} x^\alpha = 0$ for $1 \leq i \leq \ell$, where $S \subset \mathbb{N}^n$ is an arbitrary finite set of exponents and $a_{i\alpha}$ are any real numbers. We let d be the maximum degree of any monomial x^α for $\alpha \in S$. We homogenize the equations with a new variable x_{n+1} . Explicitly, define $S' \subset \mathbb{N}^{n+1}$ to consist of $\alpha' = (\alpha, d - \sum_i \alpha_i)$ for all α in S and we write $a_{i\alpha'} = a_{i\alpha}$. We add a new equation with coefficients $a_{\ell+1,\alpha} = 1$ for all $\alpha \in S'$ and $b_{\ell+1} = 1$. For this system, we can clearly take $g_{1i} = 1$ and $d_1 = d$ to satisfy the condition on exponents. Furthermore, for any positive solution (x_1, \dots, x_n) to the original system of equations, (x'_1, \dots, x'_{n+1}) with $x'_i = x_i / (\sum_\alpha x^\alpha)^{1/d}$ and $x'_{n+1} = 1 / (\sum_\alpha x^\alpha)^{1/d}$ is a solution to the homogenized system of equations.

Next, we add a multiple of the last equation to each of the others in order to make all the coefficients positive. For each $1 \leq i \leq \ell$, choose a positive $b_i > -\min_\alpha \{a_{i\alpha} \mid \alpha \in S'\}$, and define $a'_{i\alpha} = a_{i\alpha} + b_i$. By construction, the resulting system has all positive coefficients, and since the equations are formed from the previous equations by elementary linear transformations, the set of solutions are the same. \square

The practical use of the construction in the proof of Proposition 9 is mixed. The first step, of homogenizing to deal with arbitrary sets of exponents, is a straightforward way of guaranteeing the existence of the matrix g . However, for large systems, the second step tends to produce an ill-conditioned coefficient matrix. In these cases, our algorithm converges very slowly. Nonetheless, Proposition 9 shows that in the worst case, systems satisfying our hypotheses can be as complicated as arbitrary polynomial systems.

Proposition 10. *There exist bilinear equations in $2m$ variables with $\binom{2m-2}{m-1}$ positive real solutions.*

Proof. We use a variation on the technique used to prove Proposition 9.

First, we pick $2m - 2$ generic homogeneous linear functions b_1, \dots, b_{2m-2} on m variables. By generic, we mean for any m of the b_k , the only simultaneous solution of all m linear equations is the trivial one. This genericity implies that any $m - 1$ of the b_k define a point in \mathbb{P}^{m-1} . By taking a linear changes of coordinates in each set of variables, we can assume that all of these points are positive, i.e. have a representative consisting of all positive real numbers.

Then we consider the system of equations

$$(15) \quad b_k(x_1, \dots, x_m) \cdot b_k(x_{m+1}, \dots, x_n) = 0, \text{ for } 1 \leq k \leq 2m - 2$$

$$(16) \quad (x_1 + \dots + x_m)(x_{m+1} + \dots + x_{2m}) = 1$$

$$(17) \quad x_1 + \dots + x_m = 1.$$

The equations (15) are bihomogeneous and so we can think of their solutions in $\mathbb{P}^{m-1} \times \mathbb{P}^{m-1}$. There are exactly $\binom{2m-2}{m-1}$ positive real solutions, corresponding to the subsets $A \subset [2m-2]$ of size $m-1$. For any such A , there is a unique, distinct solution satisfying $b_k(x_1, \dots, x_m) = 0$ for all k in A and $b_k(x_{m+1}, \dots, x_{2m}) = 0$ for all k not in A . By assumption, for each solution, all the coordinates can be chosen to be positive. The last two equations (16) and (17) dehomogenize the system in a way such that there are $\binom{2m-2}{m-1}$ positive real solutions. Finally, as in the last paragraph of the proof of Proposition 9, we can add multiples of (16) to the equations (15) in order to make all the coefficients positive. \square

ACKNOWLEDGMENTS

We thank Bernd Sturmfels for reading a draft of this manuscript. This work was supported by the Defense Advanced Research Projects Agency project “Microstates to Macrodynamics: A New Mathematics of Biology” and the U.S. National Science Foundation (DMS-0456960).

REFERENCES

- [1] Dan Bates and Frank Sottile. Khovanskii-Rolle continuation for real solutions. arXiv:0908.4579, 2009.
- [2] Dustin A. Cartwright, Siobhan M. Brady, David A. Orlando, Bernd Sturmfels, and Philip N. Benfey. Reconstruction spatiotemporal gene expression data from partial observations. *Bioinformatics*, 25(19):2581–2587, 2009.
- [3] J. N. Darroch and D. Ratcliff. Generalized iterative scaling for log-linear models. *Annals of Math. Stat.*, 43(5):1470–1480, 1972.
- [4] Jean B. Lasserre, Monique Laurent, and Philipp Rostalski. Semidefinite characterization and computation of real radical ideals. *Foundations of Computational Mathematics*, 8(5):607–647, 2008.
- [5] Jean B. Lasserre, Monique Laurent, and Philipp Rostalski. A prolongation-projection algorithm for computing the finite real variety of an ideal. *Theoretical Computer Science*, 410(27–29):2685–2700, 2009.
- [6] Daniel D. Lee and H. Sebastian Seung. Algorithms for non-negative matrix factorization. *Adv. Neural Info. Proc. Syst.*, 13:556–562, 2001.
- [7] Lior Pachter and Bernd Sturmfels. *Algebraic Statistics for Computational Biology*. Cambridge University Press, 2005.
- [8] Jan Verschelde. PHCPACK: A general-purpose solver for polynomial systems by homotopy continuation.

DEPT. OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, BERKELEY, CA 94720, USA
E-mail address: `dustin@math.berkeley.edu`